

TITLE OF THE INVENTION

Optical Transmission Systems, Devices, and Methods

BACKGROUND OF THE INVENTION

The present invention is directed generally to optical
5 transmission systems. More particularly, the invention
relates to wavelength allocation in multidimensional
wavelength division multiplexed ("WDM") optical transmission
systems.

Digital technology has provided electronic access to
10 vast amounts of information. The increased access has driven
demand for faster and higher capacity electronic information
processing equipment (computers) and transmission networks
and systems to link the processing equipment.

In response to this demand, communications service
15 providers have turned to optical communication systems, which
have the capability to provide substantially larger
information transmission capacities than traditional
electrical communication systems. Information can be
transported through optical systems in audio, video, data, or
20 other signal formats analogous to electrical systems.
Likewise, optical systems can be used in telephone, cable
television, LAN, WAN, and MAN systems, as well as other
communication systems.

Early optical transmission systems, known as space
25 division multiplex (SDM) systems, transmitted one information
signal using a single wavelength in separate waveguides, i.e.
fiber optic strand. The transmission capacity of optical
systems was increased by time division multiplexing (TDM)
multiple low bit rate, information signals into a higher bit
30 rate signal that can be transported on a single optical
wavelength. The low bit rate information carried by the TDM
optical signal can then be separated from the higher bit rate
signal following transmission through the optical system.

The continued growth in traditional communications
35 systems and the emergence of the Internet as a means for
accessing data has further accelerated the demand for higher

capacity communications networks. Telecommunications service providers, in particular, have looked to wavelength division multiplexing (WDM) to further increase the capacity of their existing systems.

5 In WDM transmission systems, pluralities of distinct TDM or SDM information signals are carried using electromagnetic waves having different wavelengths in the optical spectrum, i.e., far-UV to far-infrared. The pluralities of information carrying wavelengths are combined into a multiple wavelength
10 WDM optical signal that is transmitted in a single waveguide. In this manner, WDM systems can increase the transmission capacity of existing SDM/TDM systems by a factor equal to the number of wavelengths used in the WDM system.

Optical WDM systems are presently deployed as in point-
15 to-point WDM serial optical links ("PTP-WDM") interconnected by electrical regenerators and switches. At each regenerator in the PTP-WDM systems, the information being transmitted can be merely regenerated on the same wavelength and retransmitted through the next link or electrically switched
20 to one of a plurality of links, different fiber, and/or a different wavelength. Various electrical switch devices can be used to switch the information between the different links at each regeneration site.

As would be expected, the cost of performing optical-
25 electrical-optical conversions in PTP-WDM systems becomes extremely expensive merely to route traffic through a network. The cost of electrical regeneration/switching in WDM systems will only continue to grow with WDM systems having increasing number of optical signal channels, or wavelengths.
30 As such, there is a desire to eliminate unnecessary, and costly, electrical switching of information being transported in optical systems.

Numerous optical cross-connect switches have been proposed as alternatives to electrical switching. For
35 example, U.S. Patent Nos. 4,821,255, 5,446,809, 5,627,925 disclose various optical switch devices. A difficulty with optical cross-connect switches is that the switches become

overly complex as the number of optical channels and input/output ports on the device is increased.

As the need for high capacity WDM systems continues to grow, it will become increasingly beneficial to provide all-optical networks that eliminate the need for electrical conversion to perform signal routing and grooming in the networks. The development of multi-dimensional all-optical networks will provide the cost and performance characteristics required to further development of high capacity, more versatile, longer distance communication systems.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings for the purpose of illustrating embodiments only and not for purposes of limiting the same, wherein:

Fig. 1 shows a system or network of the present invention;

Figs. 2-5 depict optical communication systems of the present invention;

Figs. 6-9b depicts waveband selectors of the present invention;

Figs. 10-12 depict transient grating waveband selectors of the present invention;

Figs. 13-14 depict multi-node optical communication networks of the present invention;

Fig. 15 shows an overview block diagram of methods of the present invention;

Figs. 16-18 show routing methods according to the present invention;

Fig. 19 shows an example of a method for selecting a regeneration site;

Figs. 20-23 show methods of forming spectral groups;

Fig. 24 shows a method of assigning spectral group numbers; and

Figs 25a-25d show methods of combining spectral groups.

DESCRIPTION OF THE INVENTION

Fig. 1 illustrates an embodiment of an optical network or system 10 including a plurality of optical nodes 14
5 connected by optical communication paths 26. The system 10 can employ one or more transmission schemes, such as space division multiplexing, time division multiplexing, wavelength division multiplexing, etc. The system 10 will be described as being "all-optical", although advantages of the present
10 invention can be realized with other than all-optical systems.

All-optical systems are those in which optical communication paths 26 are uninterrupted throughout the system 10 so that signals can travel between nodes 11 and
15 through nodes 11 without undergoing an optical-electrical-optical conversion. When signals in all-optical systems are to be regenerated, they are removed from the optical communication paths 26 of the system 10 without terminating the optical communication path 26, so that other signals that
20 do not need to be regenerated can continue without having to undergo an optical-electrical-optical conversion. Optical signals to be regenerated are removed from the optical communication path 26, converted to electrical form, regenerated, converted back into optical form, and the
25 regenerated optical signal is inserted on the optical communication path 26 of the system 10.

The optical nodes 11 can include various optical signal processing devices, such as transmitters and receivers (which can be collocated to form an optical network gateway
30 ("ONG")), optical cross-connect switches or routers ("OXC"), amplifiers, and optical add/drop multiplexers ("OADMs"). ONGs are used to terminate an optical communication path 26, as well as interface between two systems 10, convert signals from the optical domain to the electrical domain, and convert
35 signals from the electrical domain to the optical domain, whereas OXCs and OADMs allow signals to be selectively added

and/or dropped from a communication path 26 without terminating the path 26. The system 10 can also include signal processing devices in locations other than the nodes 11, such as amplifiers located between the nodes 11.

5 The optical nodes 11 are generally configured to switch, route, demultiplex, multiplex, convert, and/or terminate pluralities of optical signal channels as groups, referred to as wavebands or "spectral groups". Examples of signal processing devices which process signal channels as spectral groups are described in more detail hereinbelow with respect to Figs. 2-14, and such devices will be referred to as spectral group routers ("SGRs"). Each SGR will typically have the capacity for a plurality of spectral groups. For example, an SGR 11 may be designed to process up to one hundred or more signal channels using five, ten, twenty-five, fifty or more spectral groups, each including one or more signal channels within each spectral group.

A spectral group is one or more signal channels that share a common path or portion of the system 10, although all signal channels in a spectral group do not necessarily have to share the identical path. As a result of commonality within a spectral group, efficiencies can be realized with a system 10 in which SGRs 11 process signal channels as spectral groups (e.g., all signal channels in a spectral group are switched to the same output path, or all signal channels in a spectral group are added or dropped from the system 10), rather than as individual channels. As a result, the SGR 11 processing signal channels from an optical communication path 26 has only as many processing operations to perform as there are spectral groups, regardless of the number of signal channels being carried by that path 26.

The optical communication paths 26 can include various guided and unguided media to provide for communication between the network nodes 11 in the system 10. The optical communication paths 26 generally include one or more optical fibers forming the paths between network nodes 11. Each path 26 can carry one or more uni- or bi-directionally propagating

optical signal channels, or wavelengths. An optical link is the communication path 26 between two nodes 11.

In all-optical systems 10, signal channels or wavelengths must be allocated to each node 11, such that the signal channels used by some nodes 11 do not conflict with signal channels used by other nodes 11. When SGRs are used, the spectral groups and the signal channels contained therein, must be allocated in a manner such that they do not contend with other spectral groups and other signal channels. A particular spectral group may extend through an entire system 10, or it may extend through only a portion of a system 10. Because the number of spectral groups that can be handled by the nodes 11 is limited, it is often advantageous to organize the system 10 so as to maximize the number of channels in each spectral group. To that end, the system 10 can be organized into several sub-networks to allow convenient grouping of channels into spectral groups.

The sub-network for one spectral group may be different than a sub-network for another spectral group passing through the same portion of the system 10. However, the edges of one spectral group must interface with another spectral group, or the signals carried on that spectral group will not continue to propagate through the system 10. At the interface of these spectral groups, the signal channels which continue to propagate will enter a new spectral group. Such transitioning of channels between spectral groups is accomplished by the SGRs. Information can be carried on one signal channel in spectral group and then carried by another signal channel in the next spectral group, although such transitions are not required. Changing the channel on which the information is carried can be effected by dropping the signal channel at a node 11, performing an optical to electrical conversion, and then performing an electrical to optical conversion onto a different optical signal channel. Alternatively, signal channels can be changed optically by performing an optical wavelength conversion. Other processing, such as filtering, reshaping, and retiming of the signal may also be performed.

A spectral group can be defined in terms of the sub-network in which it exists. A spectral group is terminated at one or more nodes 11 defining the edge of, and bounding, the sub-network. Because the spectral groups are terminated at
5 the edges of the sub-network, spectral groups can be reused in other sub-networks in the system 10.

Spectral groups allow the size, cost, and complexity of network elements 11 (e.g., SGRs) to be reduced. Of course, a system 10 utilizing such network elements will typically
10 contain many spectral groups, and the manner in which those spectral groups are formed and ordered is important to the efficient and proper operation of the system 10. Those and other features must be considered when forming systems 10 and SGRs using spectral groups. Methods of forming such systems
15 10 and SGRs are described hereinbelow with respect to Figs. 15-29.

Fig. 2 shows a more detailed portion of the system 10. Generally, the optical system 10 has nodes 11 in the form of at least one optical transmitter 12 and at least one optical
20 receiver 14, as shown in Fig. 2. Each transmitter 12 is configured to transmit information via one or more information carrying wavelengths $18_{i,k}$ contained in at least one waveband $16_{1,i}$ to the receivers 14. Each receiver 14 is configured to receive the information carried via one or more
25 of the information carrying wavelengths $18_{i,k}$. As used herein, the term "information" should be broadly construed to include any type of data, instructions, or signals that can be optically transmitted.

As shown in Fig. 2, the system 10 further includes at
30 least one intermediate optical processing node 20, such as an optical switch 22. The transmitter 12 is configured to transmit an optical signal 24 containing one or more information carrying wavelengths 18_j along signal transmission waveguide, i.e., fiber, 26 to the switch 22 via
35 input port 28. The optical processing node 20 includes one or more waveband selectors, or selective element, 30 that are configured to pass and/or substantially prevent the passage

of information in wavebands 16_i to the receiver 14 via output ports 32. Because the information is being manipulated in wavebands, the individual information carrying wavelengths 18_j within the waveband 16_i do not have to be separated in individual wavelengths to be managed and processed. Also, the individual wavelengths 18_j within the waveband 16_i be varied in the system 10 without affecting the configuration of the optical processing node 20. Wavelengths 18_j in the original signal 24 but not within the waveband 16_i are prevented from passing through to the receivers 14.

In the present invention, optical signals 24 can be produced including a number of wavebands 16, each of which may contain one or more information carrying wavelengths in a continuous band of wavelengths or a plurality of wavelength bands. For example, a waveband 16 can be defined as having a continuous range of ~200 GHz containing 20 different information carrying wavelengths 18_{1-20} spaced apart on a 10 GHz grid. The bandwidth of each waveband can be uniformly or variably sized depending upon the network capacity requirements. Likewise, the bandwidth of the waveband is not restricted, but can be varied to accommodate varying numbers of wavelengths.

Generally, systems 10 of the present invention are configured so that the optical processing nodes do not separate and process individual information carrying wavelengths during transmission from the transmitter to the receiver. Instead, optical processing nodes 20 are configured to process the information in wavebands that may include any number of individual information carrying wavelengths. The processing of information in wavebands decreases the complexity involved in processing large numbers of channels, while increasing the flexibility of optical components deployed in the transmission path between transmitters and receivers. The bandwidth and number of information carrying wavelengths within a waveband in a network can be statically or dynamically allocated depending upon the information traffic flow in a given network segment.

Fig. 3 shows a more general arrangement of the system 10, which includes a plurality of transmitter 12_n optically connected via the switch 22 to a plurality of receiver 14_m . Analogous to Fig. 2, each transmitter 12_n transmits an optical signal 24_n which includes one or more wavelengths $18_{n,j}$ through a waveguide 26_n to an input port 28_n of the switch 22. It will be appreciated that each transmitter may include one or more sources to transmit and one or more wavelength signals. Likewise, each receiver may include one or more detectors for receiving the signals.

An optical distributor 34_n , such as a demultiplexer 36 and/or a splitter 38, is provided in the input port 28_n to distribute the signal 24_n to the waveband selectors $30_{n,m}$. An optical combiner 40_m , such as a wavelength division multiplexer 42 or a coupler 44, is generally included to combine the wavelengths $18_{m,k}$ in waveband $16_{m,i}$ emerging from the waveband selectors $30_{n,m}$ and provide a modified signal $24'_m$. The modified signal $24'_m$ exits the switch through the output port 32_m and passes along waveguide 26 to the receiver 14_m .

For convenience and clarity, Fig. 3 shows only a waveband selector 30 connecting input port 28_1 to output port 32_1 . However, it should be understood that the switch 22 will generally include at least one waveband selector 30 between each input port 28 and each output port 32. It is also noted that in some networks it is not necessary that corresponding input and output ports, e.g. 28_1 and 32_1 , be connected to loop a signal back to its point of transmission. In addition, reference numeral subscripts are generally not used in the remainder of the description to simplify the nomenclature.

Transmitters 12 used in the system 10 can include one or more optical emitters and sources that provide continuous wave and/or pulsed beams, such as one or more modulated lasers as is known in the art. The transmitter 12 may also include narrow band incoherent sources such as described in U.S. Patent Nos. 5,191,586 and 5,268,910 issued to Huber or

other optical sources for producing optical signals. Information can be directly or indirectly, e.g., externally, modulated, or alternatively upconverted, onto an optical wavelength, and the information itself may be a time division
5 multiplexed signal.

The transmitter 12 can also be used to provide multiple information carrying wavelengths using techniques such as described in U.S. Patent No. 5,400,166. Multiple information carrying wavelengths can be placed on a single carrier from
10 the transmitter 12 using techniques, such as subcarrier modulation (SCM). SCM techniques are described in U.S. Patent Nos. 5,101,450, 5,134,509, and 5,301,058 issued to Olshansky, 4,989,200 issued to Olshansky et al., 5,432,632 issued to Watanabe and 5,596,436 issued to Sargis et al.

15 The transmitters 12 may be coupled to an external electrical network or part of an optical-electrical-optical (O/E/O) signal regenerator within an optical network. One skilled in the art will appreciate that the selection of the transmitter 12 and the number of information carrying
20 wavelengths will depend upon the desired information transfer rates for a particular transmitter/receiver system at the respective nodes. While the present invention provides the ability to substantially upgrade the transfer rate for the node, it does not require that older, slower nodes be upgraded
25 upon implementation of the present invention.

Consistent with the discussion regarding the transmitter 12, the receiver 14 and transmission fiber 26 does not have to be upgraded to be compatible with the present invention. In the present invention, the capabilities of the receiving
30 system can be taken in account when establishing wavebands to be transmitted to a particular receiver 14.

As shown in Fig. 4, the receiver 14 will generally be used to separate the individual information carrying wavelengths $18_{i,k}$ in each waveband 16_i contained in the
35 modified signal $24'$ and convert the information to one or more electrical signals. The receiver may include a number of a wavelength filters, such as Bragg gratings or demultiplexers,

in combination with an optical to electrical converter (O/E), such as a photodiode, to provide for direct detection of the individual wavelengths. The receiver 14 may also provide for indirect detection of the individual wavelengths, such as by using coherent detector arrangements.

Referring to Fig. 5, the system 10 may include other types of intermediate processing nodes 20, such as add and/or drop devices. The other intermediate processing nodes can be employed to selectively modify the wavebands in the signal 24' and pass a further modified signal 24'' to successive switches 22 and to the receivers 14. The subsequent switches 22 between other intermediate processing nodes 20 and the receivers 14 can be used to further process the signal 24'' to produce a further modified signal 24''' which may include waveband subset 16_{i1}. The optical add and/or drop devices/ports can be embodied as a 2x2 switch that can provide for 100% programmable add/drop capability or by employing directional devices, such as couplers and/or circulators, with or without waveband selectors 30 to provide varying degrees of programmability, as will be further discussed.

The receiver 14 can also be used to further distribute the signal 24''' as a part of an O/E/O signal regenerator. One skilled in the art will appreciate that in an O/E/O regenerator the optical wavelengths received by the receiver 14 do not necessarily have to correspond to the optical wavelengths at which the information is further transmitted.

Waveband selectors 30 generally include at least one filter, gate, and/or switch configured to pass and/or substantially prevent the passage of at least one waveband 16 received from the inlet port 28 to the outlet port 32. A signal is generally considered to be substantially prevented from passage, if the signal is sufficiently attenuated such that a remnant of the attenuated signal that passes through the waveband selector does not destroy signals that have been selectively passed through the optical processing node 20. For example, a 40 dB attenuation of a signal will generally be sufficient to prevent cross-talk interference between

remnant signals and signals being selectively passed through the optical processing node 20.

In an embodiment shown in Fig. 6, the switch 22 includes a waveband demultiplexer 36 and an optical signal splitter 38
5 coupled via a doped optical fiber 46 to the multiplexer 42 at the output port 32. When an optical signal is to be passed to the output port 32, the doped fiber is supplied with energy from the switch pump 48 to overcome the absorption of the doped fiber 46. The amount of energy supplied by the pump 48
10 can be controlled to selectively amplify or attenuate a signal being passed through the waveband selector 30. In the absence of optical pump energy, the doped fiber 46 will absorb the optical signal, thereby substantially preventing the passage of that portion of the signal to the outlet port 32. In the
15 embodiment of Fig. 6, the wavebands can be switched to any number of output ports including one to one switching and one to many broadcasting.

The dopant in the doped optical fiber 46 can be erbium or any other dopant including other rare earth elements that can
20 render the fiber transmissive in one state and substantially less transmissive in another state. The selection of a dopant in the doped fiber will depend upon the information carrying wavelengths that are to be switched in the system. Also, mechanical, electro-optic, liquid crystal, semiconductor, and
25 other types of switches along with gratings, filters and gates, can be substituted for or used in combination with doped fiber 46 to achieve desired characteristics in the switch 22.

The waveband selector 30 may include reflective ($\geq 50\%$
30 reflectance) and/or transmissive ($\leq 50\%$ reflectance) selective elements that can be used to pass, either reflect or transmit, any of the wavebands 16 that comprise the signal 24. The waveband selector 30 may employ Mach-Zehnder filters, Fabry-Perot filters, and Bragg gratings to perform the waveband
35 selection.

As shown in Figs. 7 and 8, waveband selectors 130 and 230, respectively, can include a plurality of in-fiber

reflective Bragg gratings 50 (Fig. 6) and/or transmissive Bragg gratings 52 (Fig. 7) to pass selected wavebands to the output ports 32. Each grating, 50 and 52, can be provided to pass selected wavebands to output ports 32. Alternatively, 5 the waveband selector 30 may include a series of multiple Bragg gratings that provide for piecewise coverage of the waveband. In the case of a multiple grating waveband selector, some separation of the wavelengths in the waveband will occur between gratings, but the multiple gratings are 10 collectively operated to pass or substantially prevent the passage of the waveband. The multiple grating selector can be tuned to individual idler gaps or telescoped to one or more common idler gaps to decrease the idler gap bandwidth.

The number of gratings in Figs. 7 and 8 is shown as being 15 equal to the number of wavebands 16 being switched. However, the number of selectors provided in the switch does not necessarily have to correspond to number of wavebands 16 currently in the system. For example, the configurations shown in Figs. 6-12 may also be suitable for use in add/drop 20 multiplexers, as well as demultiplexers or multiplexers, in which any number of wavebands can be processed.

It may also be advantageous to provide sub-wavebands within the wavebands 16 of varying size that can be received, divided into the sub-wavebands, and the sub-wavebands can 25 further transmitted to other receivers within the system. The waveband selectors 30 can also be used to pass multiple wavebands to reduce the number of components in the system 10. In addition, the wavebands 16 can be selected to overlap to allow one or more wavelengths 18 to be transmitted in multiple 30 wavebands 16.

As shown in Fig. 7a, a waveband selector 130 can include a three port circulator 54 used in conjunction with the plurality of reflective Bragg gratings 50 using a configuration similar those discussed in U.S. Patent Nos. 35 5,283,686 and 5,579,143 issued to Huber, and 5,608,825 issued to Ip. In Fig. 7b, a waveband selector 230 employs transmissive gratings 52 to transmit selected wavebands to the

output ports 32 and reflect the remaining wavebands. An optical isolator 56 can be incorporated to prevent reflected wavebands from propagating back to the input ports 28. One skilled in the art will appreciate that directional couplers and other directional devices can be substituted for the optical circulators with appropriate circuit modifications.

The optical processing node 20 may include a wavelength converter 58 to provide for switching one of more of the wavelengths in the transmitted signal 24. In Fig. 7a, the wavelength converter 58 is shown before the waveband selector 30; however, the wavelength converter 58 may also be positioned after the waveband selector 30 and operated accordingly.

Similarly in Fig.8, a waveband selector 330 can be used with one or more directional devices, such as a circulator or a coupler, with either reflective or transmissive waveband gratings, 50_i or 52_i, to select wavebands. It will be appreciated that the selector 330 can be employed as an add and/or drop device/port, as well as a filter or in a demultiplexer or multiplexer in the system 10.

The optical distributor 34 associated with the input port 28 can be embodied as an optical splitter to split the signal 24 and distribute a portion of the entire signal 24 to each of the output ports 32. As shown in Fig. 9a, the optical distributor 34 can be embodied as a circulator 54 to provide the entire signal to each waveband selector 430. Wavelengths within waveband of the selector 230 are transmitted to the output port 32, while the remaining wavelengths are reflected by the transmissive gratings and circulated to successive ports.

Likewise, optical couplers can serve as the distributor 34 to provide the entire signal to waveband selector 530 (Fig. 9b). One skilled in the art will appreciate that directional devices, such as multiple three port circulators and/or coupler, can be cascaded in various other configurations equivalent to those shown in Figs. 9a&b. The gratings, 50 or 52, could be prepared having a reflectivity and transmittance

of less than 100%, to allow a portion of signal to be transmitted and reflected.

The fiber Bragg gratings 50 and 52 used in the switch 22 can be permanently and/or transiently produced. Embodiments of the present invention incorporate fixed and/or tunable permanent Bragg gratings, 50 and 52 as the waveband selectors 30. The permanent gratings used in the present invention can be prepared by conventional methods, such as by using ultraviolet (UV) light to irradiate a GeO₂ doped fiber core. Such methods are discussed in U.S. Patent Nos. 4,725,110 issued to Glenn et al., 5,218,655 and 5,636,304 issued to Mizrahi et al., which are incorporated herein by reference, and related patents.

The permanent gratings can be tuned to provide for reflectance of a waveband in one mode and transmittance in another mode. Tuning of the grating properties can be accomplished mechanically (stretching), thermally, or optically, such as discussed in U.S. Patent Nos. 5,007,705, 5,159,601, and 5,579,143, and by M. Janos et al., Electronics Letters, v32, n3, pp. 245-6, electronically, or in some other appropriate manner.

A limitation of tunable permanent gratings is that a portion of the wavelength band can not be used to transfer signals. The unused portion of the wavelength band, called an "idler" gap, is necessary to provide each permanent grating with a gap in the wavelength spectrum in which the grating will not affect a signal encountering the grating.

Transient reflective or transmissive gratings, 50^T and 52^T, respectively, could also be used in the waveband selector 30. Transient gratings can be used to reduce or eliminate the need for idler gaps in the transmission wavelengths and provide increased flexibility in the wavelength selectivity of the switch 22.

Transient gratings, either 50^T or 52^T, can be formed in a portion of the fiber in which the refractive index of the fiber can be transiently varied to produce a grating. In an embodiment, the fiber portion is doped with Erbium, other

rare earth elements, such as Yb and Pr, and/or other dopants that can be used to vary the refractive index of the fiber to produce a grating. In another embodiment, the transient grating can be formed in a fiber section that contains a
5 permanent grating to provide a combined performance grating and/or to establish a default grating in the absence of the transient grating.

As shown in Figs. 10-12, transient gratings can be written by introducing a grating writing beam either directly
10 into the transmission fiber or by coupling the writing beam into the transmission fiber. One or more transient grating writing lasers 60_i are used to introduce a transient grating writing beam into the doped portion of the signal waveguide 26. In a waveband selector 630 shown in Fig. 10, the writing
15 beam is split into two paths and introduced into the transmission fiber 26 via ports 62. A plurality of narrow wavelength reflective gratings 64_i are positioned in one of the writing beam paths to control the position of the standing wave in the waveguide 26 by introducing a time delay
20 on the wavelengths of the writing beam. Narrow wavelength reflective or transmissive gratings, 64_i or 66_i , can also be used to remove the writing beam from the transmission fiber 26.

As shown in Fig. 11, the writing beam can also be
25 reflected back upon itself using spaced narrow wavelength reflective gratings 64_i , to form a standing wave and produce a transient gratings 50^T in waveband selector 730. The grating writing lasers 60_i can be operated in conjunction with modulators 68 and pulsing switches 70 to control the
30 coherence of the writing lasers 60_i and the resulting transient gratings 50^T_i . A waveband selector 830, shown in Fig. 12, can also be configured with a reflector 72 in a coupled fiber to establish a standing wave by reflecting the writing beam back upon itself to form the standing wave in a
35 manner similar to that described with respect to Fig. 11.

Single wavelength continuous writing beam arrangements have been used for signal identification and pattern

recognition as discussed by Wey et al., "Fiber Devices for
Signal Processing", 1997 Conference on Lasers and Electro-
Optics, Baltimore, MD. Also, U.S. Patent No. 5,218,651
issued to Faco et al., which is incorporated herein by
5 reference, describes two beam methods for producing a
transient Bragg grating in a fiber.

In systems 10 of the present invention, the switch 22
can be used to optically connect a transmitter and a receiver
(Fig. 2) in a 1x1 configuration or a plurality of nodes 100
10 in an nxm configuration (Figs. 13-14). In a 1x1
configuration, the switch 22 can be useful for dropping
wavebands or for varying the waveband characteristics (gain
trimming) of the signal.

The nodes 100 used in the system 10 may contain various
15 system components including optical transmitters, receivers,
and/or other processing equipment, such as switches depending
upon whether the node is an origination (transmitting
signals) and/or a destination (receiving signals) node, and
whether it is a terminal node. The system 10 may further
20 include other optical transmission equipment, such as optical
amplifiers 74, and other optical processing nodes 20, such as
optical add/drop multiplexers, between the switches and the
nodes 100 as may be useful in a given system.

The 4x4 switch arrangement shown in Fig. 13 is
25 representative of a north-south-east-west communication
system. One skilled in the art will appreciate that the
nodes/switch arrangements can be varied to accommodate
various network configurations. For example, a 3x3
arrangement is shown in Fig. 14. The arrangements in Figs.
30 13 and 14 show the cross connections of the switches 22, but
do not show the waveband selectors within the switches 22.

The flow of communication traffic between the nodes can
take place using a variety of optical waveband hierarchies.
In an embodiment, the optical wavebands are established and
35 wavelengths assigned based on both the signal origination
node and the signal destination node to avoid the need for
wavelength conversion in the optical network.

For example, the spectrum of wavelengths used with each receiver can be divided into wavebands and the destination wavebands assigned to transmitters. The assignment may be static or dynamically controlled at the network management level so no overlap occurs in the wavebands assigned to each transmitter from the various receivers. Dynamic control of the waveband assignment provides flexibility in the wavelength management in the system and can be performed at various points in the system, such as at the client system, e.g., SONET, SDH, ATM, IP interface with the optical network.

Waveband hierarchies in which the origination and destination nodes are paired are particularly useful in communication systems in which a signal is being sent from the origin to one destination, such as in telephone communication systems. In addition, the present invention can also accommodate the necessary protection systems to provide multiple paths to the same destination by proper allocation of the wavelengths.

In a multiple destination system, such as a cable television system, it may be more appropriate for the wavebands to be determined based solely on the origination node of the signal. Waveband selectors can be included in the switches to pass signals corresponding to a particular source to any number of destination nodes. The switch can provide further control over the distribution of signals by passing broadcast signals to a distribution segment only upon a subscriber's request. The CATV provider, in response to a programming request, can centrally control the switch to deliver the signal to the requester. In the absence of an express request by a subscriber the signal would not be broadcast to the segment. The limited availability of the signal on a segment may discourage pirating of programming signals.

Switches of the present invention can also be used for remote switching and routing of communication traffic in the event of a fault in the system. For example, in Fig. 13

if a signal were to travel from node A to node C, the typical path would be through the switch connected between nodes A and C. However, if a fault occurs in the line from the switch to node C, it may be desirable to route traffic from node A through node D to node C. Upon detection of the fault, the network management system could reconfigure the switches 22 in the system 10 to reroute the traffic or switch to a previously configured protection route.

As can be seen, the present invention provides for flexibility in optical transmission systems. In addition, the present invention provides for increased transmission capacity without the commensurate increase in complexity that was present in the prior art systems.

Fig. 15 is a block diagram of methods that can be used to form systems 10 and SGRs 11 using spectral groups according to the present invention. Those methods allow for reducing the amount of regeneration, efficiently utilizing bandwidth, and providing a high degree of reliability. Those methods may be performed individually or in combination with all or some of the other methods, or in conjunction with other methods, such as span engineering. Those methods include routing connections 620, selecting regeneration sites 622, forming spectral groups 624, assigning spectral group numbers 626, and assigning channels 628. Each method will be discussed hereinbelow. Advantages of the methods can be realized with all-optical systems as well as optical/electrical systems, although most of the discussion will be focused on all-optical systems 10.

The method of routing connections 620 will be described with respect to Figs. 16-18. The method includes determining or specifying the topology of the system 10, including the endpoints of each link and the link distance. The method also includes determining or specifying a set of demands, such as the data rate and the desired level of protection. The method also includes determining or specifying other parameters, including span loss, fiber type, etc.

For all-optical systems 10, the method 620 includes shortest path routing as the starting point, as will be further explored hereinbelow, because costs generally increase with the number of regeneration sites. In contrast, 5 the method includes minimum hop routing for systems where regeneration occurs at every node, because minimizing the number of hops minimizes the number of nodes that are crossed and, therefore, the number of regenerations. Advantages of the present invention may be realized with both shortest path 10 routing and minimum hop routing, and combinations thereof.

The routing method 620 may be performed for both unprotected demands and protected demands, and includes determining the protection level on a per-demand basis or on a group or some other basis. For unprotected demands, routing 15 includes using a shortest path method. For 1+1 protection, shortest dual-path routing can be used, where, if possible, the resulting primary and secondary paths are node and link disjoint. In general, the shorter of the two paths is taken to be the primary path.

20 The routing method 620 includes forming dual paths that are not the shortest path plus the second shortest path. As shown in the hypothetical network of Figs. 16a and 16b, the shortest path between Omaha and Los Angeles is shown by the thick lines. However, selecting this path precludes a second 25 path that is completely link and node disjoint. The optimal dual-path is shown in Fig. 22.

The routing method 620 includes a mode in which 'shortest' is the shortest distance. In this mode, routes are chosen such that the end-to-end path (or the combination of the 30 primary and secondary end-to-end paths) has minimum distance. In another mode, the method includes factors such as regeneration sites, minimizing regeneration sites, optical equipment penalties, and noise figures, so that routes are defined using factors such as the Optical Signal to Noise 35 Ratio (OSNR), which is associated with the Noise Figure associated with each link. In that mode of operation, routing includes 'minimum noise figure' routing.

Another factor that can be used is "optical penalty". Each piece of optical networking equipment (e.g., Router, OADM, etc) has an optical penalty associated with it. When calculating shortest path (whether based on distance or noise figure), the routing method 620 includes determining the optical penalties of routes. For example, the method can calculate an effective distance for each link that is comprised of the actual distance of the link, plus half of the optical penalty at one link endpoint, plus half of the optical penalty at the other endpoint. Using half of the penalty at each end allows the link direction to be ignored when calculating shortest path. If a node is an intermediate node on a path, the path will contain a link into and out of the node, thus, the whole optical penalty will be counted at that node. For end points half of the optical penalties can be counted, although when determining the shortest path between a set of endpoints, the penalties at the endpoints do not play a role since they will be the same for all paths.

Furthermore, the method 620 includes distinguishing between the penalty suffered by traffic that optically passes through a node 11 as opposed to traffic adding/dropping at the node 11, and assigning different optical penalties for the different traffic patterns. For example, optically passing through an OADM will result in a noise figure and distance penalty. If signal channels are added/dropped at an OADM, the corresponding penalties may be different.

In addition to determining the shortest path, the method 620 also includes selecting one or more possible alternative paths for each connection. The method 620 includes generating a list of possible paths. For example, the method 620 includes routing all connections along their most desirable path (without regard to link capacity), determining where the hot-spots are in the network (i.e., the spots in the network that are likely to become congested), searching for alternative paths that avoid these hot spots (while enforcing additional rules such as alternative paths should have no more than N extra regenerations - we have found that N can

often be set to 0 and still provide useful alternative paths), and identifying one or more of the alternative paths. In some cases, an alternative path may have fewer regenerations than the "shortest" path (e.g., due to
5 electronic termination sites being present, or due to the actual lengths of the primary and secondary), in which case the method includes selecting the alternative path as the preferred path.

After generating the shortest path and any alternative
10 paths for each connection, routing of the demands is performed. For each connection, the method includes choosing a path such that the network is left in the least-loaded state (i.e., the Least-Loaded algorithm). In other words, the method identifies the "most loaded link" in the path, and
15 can chose the alternate path in which the most load link carries the lightest load.

In general, it has been found that the relatively sparsely interconnected systems 10 do not result in a large number of alternative paths. Thus, alternative path routing produces
20 only small gains in these networks. It has also been found that, for purposes of efficiently packing spectral groups, routing connections over the same path, rather than splitting them up over several paths, produces somewhat better results.

As systems 10 become more heavily loaded, the shortest
25 path, or shortest dual-path, for a particular connection may not have enough free capacity to support the connection. Rather than immediately add a second fiber to the congested links, the method includes searching for alternative paths when such new demands are made. Even if the alternative path
30 requires more regeneration, it may be less expensive than populating a second path 26 with amplifiers and upgrading the nodal equipment to support the additional path 26.

In some systems 10 it is desirable to provide 1+1 protection paths, but the presence of spurs or other
35 degenerate topologies in the topology prevents completely node- and link-disjoint primary and secondary paths. The routing method includes identifying common nodes and/or

common links which prevent complete disjointer and, if appropriate, authorizing the use of common nodes or links for protection paths. In that case, the method includes finding the shortest dual path, such that the number of common
5 nodes/links is minimized. If in the routing process, the shortest such path does not have free capacity, alternative paths can be considered, with or without limitations on the number of common nodes/links. Fig. 17 illustrates an example of 1+1 protection over a common link and nodes (i.e., nodes D
10 and E and the link therebetween). In that example, the degenerate case is not a spur, although a system 10 having a spur is analogous.

The method of the present invention can allow for common nodes and common links to be authorized. As another example,
15 in the network of Fig. 18, if the method allows for common nodes but not common links, a primary and secondary path can be found from A to E using common node B (i.e., paths ABCE and ABDE), but not from A to F, which would require the authorization of link EF to be in common.

The method for selecting regeneration sites 622 will be described with respect to Fig. 19. The method 622 is particularly relevant in all-optical systems 10. The method
20 622 can be used in conjunction with one or more other methods, such as to form all-optical systems 10, or it may be used alone to plan regeneration sites. The method 622 has
25 several features, one or more of which may be used.

The need to regenerate channels, or optical reach, in a spectral group depends on the maximum distance that the channels can travel without loss of data. The optical reach
30 can vary from channel to channel depending on such factors as fiber type, amplifier span lengths, optical equipment penalties, etc. There is often flexibility as to where regeneration can occur, as illustrated in the example of Fig. 19. Assume that the connection from A to D needs to be
35 regenerated once, and that regeneration can occur at either node B or C. In general, spectral group packing is more efficient when regenerations occur at the same site. Thus,

the method can be made to choose certain sites for regeneration. For example, node B may be consistently chosen over node C as the regeneration site. This is not to say that the site becomes a dedicated regeneration site. Connections not requiring regeneration will still optically bypass the site.

Also, the method can specify some sites as 'not preferred for regeneration' (e.g., an office with limited space). At such sites, regeneration will only occur if absolutely necessary. Thus, in the example of Fig. 19, if B is a non-preferred regeneration site, node C would be chosen instead.

The method can determine whether the system should include an optical/electrical hybrid portion including dedicated regeneration sites. For example, if a link has a distance of 6,000 km, and if the system reach is 4,000 km, a dedicated regeneration site could be added along this link in an appropriate location. The site could be, for example, 'back-to-back' transmitters and receivers.

The regeneration method can also select regeneration sites to vary the amount of equipment required in the system 10. Selective regeneration of the signal channels in one or more spectral group can be performed at a node 11 to avoid the need to deploy a regeneration node where such a node would not otherwise be needed.

Regeneration site selection can also take into account other factors, such as the distance, accumulated noise figure, OSNR, etc. For example, if the method is in the 'route on distance' mode, regeneration is based on the length of the links and the 'distance penalties' of the equipment.

Another factor is the maximum number of optical nodes 11 that can be transited before transients and/or accumulated crosstalk becomes a problem. If a threshold is crossed, the signal is regenerated, independent of the length of the path.

The method of forming spectral groups 624 will be described with respect to Figs. 20-23. As with the other methods, the method of forming spectral groups 624 may be

used in conjunction with one or more other methods, or it may be utilized by itself. If the method is used to form a system 10 using spectral groups, it can be advantageously utilized after the methods of routing connections 620 and selecting 5 regeneration sites 622.

Spectral groups typically are designed to have a maximum number of channels, and can be populated with any number of channels from zero to the maximum.

The method of forming spectral groups 624 includes 10 determining whether the segment of the connection that falls between two regeneration points (or between an endpoint and a regeneration point) is a subconnection (if a connection has no regenerations, then the subconnection is equivalent to the connection). In general, spectral groups are formed by 15 grouping subconnections, not end-to-end connections.

Initially, all subconnections that have the same two endpoints are bundled together. It is usually desirable to minimize the number of bundles formed. An important feature of the method is grouping the channels in such a way as to 20 form spectral groups with a high fill rate. However, there are situations where a bundle must be split into multiple smaller bundles (e.g., moving some subconnections to a new bundle) For example, at SGR nodes it is advantageous to bundle together subconnections that will be transmitted and 25 received at the same location. The subconnection bundle can be split such that all subconnections in the resulting smaller bundles have the same transmitter and receiver location. Also, if a bundle contains more subconnections than 30 the maximum number of channels in the spectral group, it needs to be split such that the resulting smaller bundles contain no more than this maximum.

After the bundles are formed they are referred to as spectral groups (SGs). The method includes determining how many spectral groups have been formed on each link. The 35 maximum number of SGs is a feature of the system 10 and the nodes 11 which is determined or provided earlier in the method.

If the maximum number of SGs is exceeded on a particular link, then the spectral grouping method includes combining or joining two or more SGs together to form a single SG. Three processes will be discussed: subsetting, merging, and
5 branching. Those processes can make use of consolidation and/or broadcast features in various optical processing devices to reduce the number of spectral groups used in the optical link. However, such reduction typically results in some unusable bandwidth, as will be further described. One or
10 more of the combining methods can be used in allocating connections to spectral groups and channels.

The steps of combining SGs include checking that the constituent subconnections can be combined into the same SG. For example, the total number of subconnections assigned to
15 channels in the resulting SG cannot be greater than the maximum number of channels allowed. Other channel allocation rules, can also be included in the combining steps. Furthermore, the combining steps can include determining whether it is possible to combine certain SGs because of the
20 need to keep the subconnections corresponding to certain primary and secondary connections in separate SGs.

The first SG combining method is subsetting. Consider the example shown in Fig. 20, which illustrates a SG that runs from Node 1 to Node 4 that contains two subconnections, a
25 second SG that runs from Node 1 to Node 3 that contains one subconnection, and a third SG that runs from Node 2 to Node 4 that contains one subconnection. All three SGs can be combined to form a single SG that extends from Node 1 to Node 4, with the resulting SG containing four channels, each of
30 which is assigned to a subconnection.

The example of Fig. 20 illustrates the broadcast, or 'drop and continue', feature of the system. Individual wavelengths can be dropped from the SG at any node. It is not necessary that all wavelengths in the SG have the same source and destination nodes. In Fig. 20, for example, the SG drops
35 channel 3 at Node 3, but the SG continues on to Node 4 where the spectral group is terminated.

It is noted that in the Fig. 20 example, channel 3 is unavailable for use on the link between Nodes 3 and 4, and channel 5 cannot be used between Nodes 1 and 2.

The subsetting method can consider efficiency when performing the subsetting operations. For example, a spectral group that extends over M links has a capacity of M*N channel slots, where N is the maximum number of channels per spectral group. The efficiency, or fill-rate, is determined by dividing the actual number of channel slots used by the maximum capacity. In the example shown in Fig. 21, the total capacity is 3*8 =24 channel slots. Of these, 18 are used, yielding a fill-rate of 75%. Two factors contribute to the fill-rate: the amount of unusable bandwidth (for example, Channel 5 from Node 3 to Node 4), and the amount of unused bandwidth (for example, in the figure, there are only seven subconnections in the SG). In lightly loaded networks fill-rate can be normalized based on the number of channels used in the SG.

Accordingly, the subset method can first join together subconnections to increase the fill rate where the resulting fill rate exceeds a certain threshold, or efficiency. If that is still not sufficient to produce a feasible solution, the method can reiterate, with the efficiency threshold lowered.

The merging method operates to merge channels in two SGs from different links that are the same or include at least one common channel. In the example of Fig. 22, one SG extends from Node 1 to Node 4, and a second SG extends from Node 2 to Node 5, with both SGs containing two subconnections. The two SGs can be merged into a single SG that extends from Node 1 to Node 5. The method can check that the path from Node 1 to Node 5 does not require regeneration due to the accumulation of noise from the merged spectral groups. The method can also check that no optical rings have been formed when merging the SGs.

In the example of Fig. 22, channels 1 and 2 cannot be used on the link between Nodes 4 and 5, and channels 3 and 4 cannot be used on the link between Nodes 1 and 2. As with

subsetting, the percentage of overlapping links can be taken into account when the merge method is performed.

The branching method can also be used to combine spectral groups. While subsetting and merging take advantage of the ability to drop one or more channels at a node and continue the SG to further nodes, the branching method takes advantage of broadcasting over multiple links at a node (such as a node equipped with an SGR). Consider the example of Fig. 23. One SG runs from Node 1 to Node 5, another from Node 1 to Node 6. The two SGs can be combined to form a single SG. In the West to East direction, the optical router at Node 4 broadcasts the SG onto the link between Nodes 4 and 5 and the link between Nodes 4 and 6. In the reverse direction, the optical router consolidates the signals from these two links. Note that it would also be possible to have a third subconnection in this SG that extends from Node 1 to Node 4, whereby the optical router would be performing a 3-way broadcast and consolidation.

In the example shown in Fig. 23, channel 1 is not utilized on the link between Nodes 4 and 6, and Channel 3 is not utilized on the link between Nodes 4 and 5. As with the other combining operations, the percentage of overlapping links between SGs is considered.

The branching method tends to be more useful in a highly interconnected system than in a sparsely interconnected system. As with merging, the noise accumulates with each link. Thus, in the example of Fig. 23, if the branching method calculates the OSNR, it should do so over all five links to determine if the resulting SG is within the regeneration bounds.

The method of forming spectral groups will be described with respect to Figs. 20-23. The regeneration method may be used independently or it may be used in conjunction with one or more other methods. For example, it may be performed prior to the routing step for forming spectral groups. Alternatively, the regeneration method may be used alone to plan regeneration sites. The regeneration

method has several features, one or more of which may be used.

The method of forming spectral groups 624 can be used to fit new subconnections into existing SGs when new demands are added to an existing network. For example, the method can fit subconnections into SGs that have the same source and destination, analogous to the initial steps of the spectral grouping method. For example, the first steps can be fitting the new subconnections into spectral groups including the same subconnections. If the number of resulting SGs (existing plus new ones) is too high, the combining method goes into the combination mode, where it performs subsetting, merging, and branching, as described above. The method can initially move new subconnections into existing SGs as opposed to combining new SGs. The percentage of overlapping links is considered when the new subconnections are combined with existing or new SGs.

The method of forming spectral groups 624 can extend an existing SG. For example, an existing SG may extend between Nodes 1 and 2. A new subconnection that extends from Node 1 through Node 2 to Node 3 may be combined with this existing SG. Thus, the existing SG would be extended to include the link from Node 2 to Node 3. The switching process involved in extending the SG should not affect existing in-service subconnections.

The method of assigning spectral group numbers 626 will be discussed with respect to Fig. 24. This method 626 can be used alone or in combination with other methods. For example, it can be used after the spectral group formation method 624. The spectral group assignment method assigns each spectral group a number from 1 to the maximum number of SGs. The SG numbers can be assigned such that two SGs are not assigned the same SG on a given fiber on a given link, or to any transmitter/receiver. In Fig. 24, four SGs are shown. The two SGs that ride on the link between Nodes 1 and 2 cannot be assigned the same number (assuming there is just one fiber pair per link); similarly, the two SGs that drop at the same

transmitter/receiver of Node 3 cannot be assigned the same SG even though they have no links in common.

In addition to selecting an SG number, the method can select a fiber pair for each link of the SG (this is relevant for links with multiple fiber pairs). The method can also make other selections, such as selecting a transmitter/receiver or other device in situations when such selections are available.

The method of assigning SG numbers includes ordering the SGs based on the difficulty of finding a satisfactory assignment. The number of SGs that have been formed on each link and on each optical router ONG is tallied. Each SG is assigned a weighting based on its most heavily utilized link or nodes. Thus, SGs on links or nodes with the most SGs will be assigned first so that there is the most flexibility in assigning the SG number (i.e., as SGs are assigned numbers, there is less flexibility in choosing numbers for the remaining SGs since they have to avoid conflict with those already assigned). For SGs on lightly loaded links or nodes, their assignment order can be based on the number of hops comprising the SG. The greater the number of hops, the more difficult it is to find a free SG number on each one of the hops.

The method can use the Most Used scheme, where an SG number is considered for assignment based on the number of times it has been assigned already. The rationale for this strategy is that the more times a particular SG number is used, the harder it is to assign it to another SG, because the chance of conflict is higher. Thus, if it can be assigned without conflict, it is chosen. If the Most Used scheme does not successfully find an SG assignment that satisfies all SGs, a Least Loaded scheme can be used. The least loaded links and nodes are chosen at each step if possible. In most scenarios there is only one fiber pair per link so that this scheme essentially operates as a *First Fit* scheme that takes into account loading at the node.

The method of assigning spectral group numbers can also consider factors such as certain SG numbers have restricted optical reach on certain fiber types. This can be taken into account when SG numbers are assigned. When the method
5 utilizes that factor, it includes the step of determining the fiber types comprising each SG and may include further steps to determine whether the fiber types vary on a span-by-span basis. Such assignment schemes include considering the SG numbers with the smallest reach that are still greater than
10 the length of the SG.

Another consideration when assigning SG numbers is that on lightly loaded systems, it is beneficial to have active wavelengths that span the full spectrum, so that the amplifier gain flattening algorithms work properly. Thus,
15 initially, SG numbers can be assigned so that such a spread is achieved, if possible (e.g., SGs may be assigned in the order 11, 5, 17, not simply 1, 2, 3).

Wavelength inventory can also be considered when assigning spectral group numbers. For example, customers may
20 have certain Tx/Rx wavelengths in inventory that they wish to use. The SGs chosen based on the methods described above may not match up with these wavelengths. The spectral group method, however, can utilize the unused Tx/Rx modules that are located at each node and any Tx/Rx modules that are
25 located in a central inventory location. The method can chose to match up with the Tx/Rx in inventory as opposed to the SG numbers that would otherwise be chosen. Preference can be given to using the Tx/Rx that are located at a particular node rather than the ones located in a central 'bin.' This
30 mode of operation may lead to a small decrease in the utilization of the network because the SG assignment process may be less than optimal.

In general, the methods can be operated to avoid single points of failure when implementing protections schemes. For
35 example, avoiding common nodes and links, avoiding common components within a node or, if that is not an option, avoiding common cards within a common component.

Even if the number of SGs formed on each link and ONG is less than or equal to the maximum allowable number, there is no guarantee that the above SG assignment processes will find a feasible solution. There may be too many constraints such
5 that a solution cannot be found. In this case, the SG bundling process can reiterate with a lower efficiency threshold for combining two SGs into a single SG.

If further iterations of the methods still does not produce a feasible solution, then the method can return to
10 the routing step. For example, the method can determines which link has the most SGs, and reduces the maximum number of paths that can be routed on that link. The whole process of routing, regenerating, SG bundling, and SG assignment can be performed again. The process can continue to iterate until
15 a solution can be found.

The method of assigning channels will be described with respect to Fig. 20 and can be performed individually or in combination with other methods. If the channel assignment method is performed with other methods to form spectral
20 groups, it can be performed after the SGs have been assigned numbers.

The methods can operate so that SG and channel assignment is performed on a subconnection basis, so that there is no attempt to assign the same SG number to all subconnections
25 that comprise a connection. Thus, whenever regeneration occurs (which partitions a connection into subconnections), wavelength conversion typically occurs. Sparse wavelength conversion can be very effective in yielding high utilization in a system 10.

Figs. 25a-25d are flow charts illustrating one
30 embodiment of the method of forming spectral groups 624 described above. In Fig. 25a the method begins by routing a connection between an origination node and a destination node. Routing the connection can be subjected to certain
35 criteria. For example, if no protection paths are required, routing a connection may include finding the "shortest" route between the origination node 11 and the destination node. The

"shortest" path is not necessarily the path with the shortest fiber length. For example, determining the shortest path may include considering factors such as fiber type, fiber condition, equipment condition, link distances, signal
5 regeneration, signal degradation, amplifier types, and other considerations and optical penalties of a particular path. If a protection path is required, then two paths must be found. For example, in a 1+1 protection scheme, the two paths are node disjoint (except at the origination and destination
10 nodes) and thus link disjoint, which may result in the shortest protection path being longer than the shortest overall path. Nonetheless, the shortest path that satisfies the protection criteria and the second shortest path that satisfies that criteria are typically the primary and
15 secondary paths, respectively. At the origination and destination nodes 11, which are shared even in 1+1 protection schemes, criteria may include ensuring that the primary and secondary paths are assigned to different transmitter/receiver equipment and routers so as to avoid
20 single points of failure in the origination and destination nodes 11.

Once the connections are made, they can be partitioned into smaller sub-connections or sub-networks, such as various optical links interconnecting nodes 11 in the system 10,
25 which is often desirable in large or long-haul systems 10. Thereafter, the connections and/or sub-connections that share the same paths are identified and bundled together to form spectral groups and sub-networks. One or more spectral groups are assigned to those bundled connections and/or sub-
30 connections which share the same paths. The reason more than one spectral group may be assigned is that the number of channels may exceed the capacity of a single spectral group. If the number of assigned spectral groups is less than a maximum number allowed number of spectral groups, then
35 individual signal channels within the assigned spectral group can be assigned to each information channel to provision the system 10.

In some networks, however, the number of information channels may be sufficiently large that too many spectral groups are initially formed. As a result, the spectral groups and sub-networks must be further refined to reduce the number of spectral groups. Such refining can be done using the
5 aforementioned concepts of subsetting, merging and/or branching, using various allocation criteria, as will be disclosed in more detail below with respect to Figs. 25b-25d. One criteria that may be used in the subsetting, merging, and
10 branching operations is a minimum acceptable efficiency of the refined spectral groups. In that way, the subsetting, merging, and branching operations will be forced to meet the minimum efficiency standards when refining the spectral groups. An inefficient system
15 can cause difficulties in the future when additional modifications are made to the network, such as adding capacity or changing traffic patterns.

Fig. 25b shows one embodiment of a subsetting operation according to the present invention. In that operation, each
20 spectral group is considered with respect to every other spectral group, to determined whether spectral groups can be combined. An example of a three part test for combining spectral groups is: (1) is the sub-network of one spectral group contained within a sub-network of another spectral
25 group, and (2) is the efficiency of the combined spectral groups greater than or equal to the efficiency level set for the system
30 10, and (3) is the total number of channels in the combined spectral group less than or equal to the capacity of the spectral group. If the test for combining the spectral groups is satisfied, the spectral groups are combined, and if the test is not satisfied, the spectral groups are not combined. The method can be iterated until all subsetting operations meeting the test criteria have been performed, or until the number of spectral groups is within the acceptable
35 range.

Fig. 25c shows a method for performing a merging operation. That method is similar to the subsetting operation

of Fig. 25b, except that the test for merging includes determining whether two sub-networks have an overlapping portion at opposite ends of the sub-networks. In the illustrated method, the test for merging also includes
5 determining whether the merged sub-networks satisfy the efficiency criteria and whether the merged sub-networks will exceed the capacity of a single spectral group.

Fig. 25d shows a method for performing a branching operation. That method is similar to the subsetting and
10 merging operations of Figs. 25b and 25c, except that the test for branching includes determining whether the two sub-networks have an overlapping portion that is in the middle of at least one of the sub-networks. In the illustrated method, the test for branching also includes determining whether the
15 branched sub-networks satisfy the efficiency criteria and whether the branched sub-networks will exceed the capacity of a single spectral group.

After completing the subsetting, merging, and branching operations, the method can return to Fig. 25a to determine
20 whether there are still too many spectral groups. If the answer is no, the method of forming the spectral groups may terminate. If there are still too many spectral groups, the method can lower the efficiency criteria, or change some other criteria, and return to the subsetting, merging, and
25 branching operations to further refine the spectral groups and sub-networks. When returning to the subsetting, merging, and branching operations, the method may return the spectral groups to their original form (as produced by the bundling operation of Fig. 25a) and re-perform the operations with the
30 lower efficiency criteria. Alternatively, the subsetting, merging, and branching operations may be performed on the spectral groups as they exist after the last subsetting, merging, and branching operations. In the later case, however, the operations may become more complex, such as when
35 calculating the efficiency of spectral groups including two or more channels with different paths. Furthermore, by not returning the spectral groups to their original form, the

operations will lose the opportunity to make certain combinations which were prohibited by the efficiency criteria in a previous iteration.

5 The methods discussed with respect to Figs. 25a-25d are illustrative of the present invention. Variations on those methods, such as searching and comparing only some of the spectral groups for the various subsetting, merging, and branching operations, as opposed to comparing all subsets to all other subsets, can of course be utilized in according
10 with the present invention. Furthermore, although the various operations are disclosed as being used together, one or more of the operations may be used separate from the rest. Also, the order of the steps and operations can be modified and still realize the benefits of the present invention.

15 Those of ordinary skill in the art will appreciate that numerous modifications and variations that can be made to specific aspects of the present invention without departing from the scope of the present invention. It is intended that the foregoing specification and the following claims cover
20 such modifications and variations.